that is best solved by distributed machines able to effectively support terminal interactions locally and coupled to other AIM machines and facilities through network or telephone links. As new machine resources are introduced into the community, we will allocate budgeted funds with Executive Committee advice to assure effective communications links.

Resource Software

We will continue to maintain the existing system, language, and utility support software on our systems at the most current release levels, including up-to-date documentation. We also will be extending the facilities available to users where appropriate, drawing upon other community developments where possible. We rely heavily on the needs of the user community to direct system software development efforts.

Within the AIM community we expect to serve as a center for software-sharing between various distributed computing nodes. This will include contributing locally-developed programs, distributing those derived from elsewhere in the community, maintaining up-to-date information on subsystems available, and assisting in software maintenance.

Community Management

We plan to retain the current management structure that has worked so well in the past. We will continue to work closely with the management committees to recruit the additional high-quality projects which can be accommodated and to evolve resource allocation policies which appropriately reflect assigned priorities and project needs. We expect the Executive and Advisory Committees to play a continuing role in advising on priorities for facility evolution and on-going community development planning in addition to their recruitment efforts. The composition of the Executive Committee will continue to represent major user groups and medical and computer science applications areas. The Advisory Group membership spans both medical and computer science research expertise. We expect to maintain this policy.

We will continue to make information available about the various projects both inside and outside of the community and, thereby, promote the kinds of exchanges exemplified earlier and made possible by network facilities.

The AIM workshops under the Rutgers resource have served a valuable function in bringing community members and prospective users together. We will continue to support this effort. In July 1984, the AIM workshop will be hosted by Ohio State University. We will continue to assist community participation and provide a computing base for workshop demonstrations and communications. We also will assist individual projects in organizing more specialized workshops as we have done for the DENDRAL and AGE projects.

We plan to continue indefinitely our present policy of non-monetary allocation control. We recognize, of course, that this accentuates our responsibility for the careful selection of projects with high scientific and community merit.

Training and Education Plans

We have an on-going commitment, within the constraints of our staff size, to provide effective user assistance, to maintain high-quality documentation of the evolving software support on the SUMEX-AIM system, and to provide software help facilities such as the HELP and Bulletin Board systems. These latter aids are an effective way to assist resource users in keeping informed about system and community developments and solving usage problems. We plan to take an active role in encouraging the development and dissemination of community knowledge resources such as the AI Handbook, up-to-date bibliographic sources, and developing knowledge bases. Since much of our community is geographically remote from our machine, these on-line aids are indispensable for self-help. We will continue to provide on-line personal assistance to users within the capacity of available staff through the MM and TALK facilities.

We budget funds to continue the "collaborative linkage" support initiated during the first term of the SUMEX-AIM grant. These funds are allocated under Executive Committee authorization for terminal and communications support to help get new users and pilot projects started.

Core Research Plans

Several systems have been developed in recent years to serve as vehicles for knowledge engineering and research on knowledge representation and its use. Knowledge acquisition (including machine learning) and advanced architectures for AI will be the two areas of most new activity in the coming year. Research on these topics obviously must draw on on-going work in representation and control.

In particular, we will focus on

- Inductive learning of MYCIN-like rules from case data in the domain of diagnosing disorders where the chief complaint is jaundice;
- Learning from experience in domains where the means for interpreting new data are largely contained in the emerging (and thus incomplete and not wholly correct) theory;
- Learning by watching a medical expert diagnose cases presented by NEOMYCIN;
- Investigating complex signal understanding systems for ways to exploit and represent concurrency with a view toward hardware and software architectures that may be capable of several orders of magnitude improvement in performance.

I.B. Highlights

During the past year, the central SUMEX machine has continued to demonstrate its important function as a "seed" environment for new investigators who are embarking on the initial stages of AIM research efforts. SUMEX thus serves as a catalyst and proving ground for new ideas. The potential of such innovations typically needs to be demonstrated in order to provide credible proposals for independent research funding. As more mature projects increasingly turn to professional workstations for their implementation and refinement, we see SUMEX's role as a source of "seed" support for new efforts as being a particularly key element in its function.

In this section we describe several of the highlights of the last year's activities. These include some older projects that have passed important milestones, new pilot projects that have showed remarkable progress in their initial stages, and some other special activities that reflect the impact and influence that SUMEX is demonstrating in the scientific and educational communities.

I.B.1. Progress Towards a Distributed SUMEX-AIM

This past year saw several technical developments at SUMEX which further demonstrate our ability and direction towards establishing SUMEX-AIM as a true distributed resource.

The SUMEX technical staff successfully completed the establishment of a remote computing facility for the Heuristic Programming Project. This new facility, located at 701 Welch Road just off of the Stanford Campus, is connected to SUMEX-AIM via a special 'twisted-pair' ethernet, designed by Nick Veizades, our Senior Electronics Engineer. This new facility also incorporates both 3 and 10 megabit/sec ethernets. The support of these two networks, along with the special ethernet link, necessitated a great deal of work in network software to accommodate this configuration. The resulting technology provides AIM researchers on Welch Road with high speed access to the SUMEX-AIM computer resource despite their remote location. This capability will be of heightened importance when the SUMEX and ONCOCIN groups join the HPP on Welch Road in new quarters sometime during the next year.

One of the most exciting computing prospects for the coming decade is the development of professional workstations. As we have discussed in prior reports, these machines may have a profound impact on biomedicine by serving as the vehicle for the practical export of expert advice systems into the hands of physicians, chemists, biologists, engineers, or other users. SUMEX has continued its investment and research into the use of workstations for biomedical AI research, and the integration of these workstations into a reliable and robust networking environment. In addition to high speed Lisp-based scientific workstations, we believe the use of low cost workstations, which offer suitable local processing power, high resolution screens with easy to use user interfaces, and networking and communications abilities, are vital to the future of our resource.

I.B.2. New Molgen Directions

For several years, the MOLGEN project has focused on research into the applications of symbolic computation and inference to the field of molecular biology. This has taken the specific form of systems which provide assistance to the experimental scientist in various tasks, the most important of which have been the design of complex experiment plans and the analysis of nucleic acid sequences. MOLGEN is now moving into a new phase of research which explores the methodologies scientists use to modify, extend, and test theories of genetic regulation, and then to emulate that process within a computational system.

The first goal of the new work in scientific theory discovery was to study extensively an existing example of the process. Professor Charles Yanofsky's work in elucidating the structure and function of regulation in the *trp* operon of E. coli provided an excellent subject that spanned twelve years of research, dozens of collaborators, and almost one hundred research papers.

Extensive interviews have been conducted with Professor Yanofsky and many of his former students and collaborators, and there has been a thorough examination of most of the relevant research papers. This has provided the MOLGEN team with a good understanding of the three major classes of knowledge that were important in the discovery of the theory of regulation in the *trp* operon: knowledge about the relevant biological objects, knowledge about the techniques used to elicit new information, and discovery heuristics used to build new models. The major stages in the discovery process have been mapped out, and work has begun on constructing a knowledge base that will represent the state of the world at the beginning of the *trp* operon research.

I.B.3. ONCOCIN - An Oncology Chemotherapy Advisor

The ONCOCIN Project, now in its fifth year, and is one of many Stanford research programs devoted to the development of knowledge-based expert systems for application to medicine and the allied sciences. The program is designed to give advice regarding the management of patients receiving cancer chemotherapy. The central issue in this work has been to develop a program that can provide advice similar in quality to that given by human experts, and to insure that the system is easy to use and acceptable to physicians. The work seeks to improve the interactive process, both for the developer of a knowledge-based system, and for the intended end user. In addition, the ONCOCIN group has emphasized clinical implementation of the developing tool so that they can ascertain the effectiveness of the program's interactive capabilities when it is used by physicians who are caring for patients and are uninvolved in the computer-based research activity. ONCOCIN is the first AIM program to have achieved routine (albeit experimental) use by non-collaborating physicians.

ONCOCIN has been used routinely in the Stanford Oncology Clinic for almost three years. Thus, much of the emphasis of this research has been on human engineering so that the physicians will accept the program as a useful adjunct to their patient care activities. The research team has pressed their effort to adapt ONCOCIN to run on professional workstations (specifically the Xerox 1108 "Dandelion") which can eventually be dedicated to full time clinic use. In keeping with other SUMEX experiments in the use of professional workstations as vehicles for implementing medical advice systems, the ONCOCIN team envisions such machines as the model for eventual non-Stanford dissemination of this kind of technology. They have been granted supplemental funding from DRR for three years to support workstation development (along with knowledge base development). They are planning to add all of the protocols in use at the Stanford oncology clinic to ONCOCIN. Major accomplishments in the past year have included the completion of formal studies to evaluate the system's impact in the oncology clinic, the development of a protocol entry system (OPAL) for use by oncologists entering new chemotherapy information into the program, and the development of an 1108 Dandelion environment that is customized for the specialized development needs of this large multi-person project.

I.B.4. New Pilot Projects

This past year saw the addition of several new SUMEX Pilot Projects. Among them are:

PATHFINDER

THE Pathfinder project is directed by Dr. Bharat Nathwani of the Department of Anatomical Pathology, City of Hope National Medical Center, Duarte, California and Dr. Lawrence M. Fagan, Department of Medicine, Stanford University. This project addresses difficulties in the diagnosis of lymph node pathology. Five studies from cooperative oncology groups have documented that, while experts show good agreement with one another, the diagnosis made by practicing pathologists may have to be changed by expert hematopathologists in as many as 50% of the cases. Precise diagnoses are crucial for the determination of optimal treatment. To make the knowledge and diagnostic reasoning capabilities of experts available to the practicing pathologist, The PATHFINDER team has developed a pilot computer-based diagnostic advice system. The project is a collaborative effort of the City of Hope National Medical Center and the Stanford University Medical Computer Science Group. A pilot version of the program provides diagnostic advice on 45 common benign and malignant diseases of the lymph node based on 77 histologic features. The group's research plan, which led to a research proposal to the NIH that is now under consideration, is to develop a full-scale version of the computer program by substantially increasing the quantity and quality of knowledge. They will also further develop techniques for knowledge representation and manipulation appropriate to this application area. The design of the program has been strongly influenced by the INTERNIST/CADUCEUS program that has also been developed on the SUMEX resource. An eventual goal is to merge the diagnostic capabilities of PATHFINDER with a microscope automation effort that Dr. Nathwani is pursuing in collaboration with experts on image processing at Carnegie Mellon University.

Protean

The PROTEAN project involves Dr. Oleg Jardetzky of Stanford Medical School's Nuclear Magnetic Resonance Lab and Prof. Bruce Buchanan of the Computer Science Department. This project has two goals: (a) to use existing AI methods to aid in the determination of the 3-dimensional structure of proteins in solution (not from x-ray crystallizing proteins), and (b) to use protein structure determination as a test problem for experiments with the AI control structure known as the Blackboard Model.

RXDX

The RXDX project is staffed by Dr. Robert Lindsay, Dr. Michael Feinberg, and Dr. Manfred Kochen from the University of Michigan and Dr. Jon Heiser, of the Metropolitan State Hospital in Norwalk, California. This project is developing a prototype expert system to act as a consultant in the diagnosis and management of depression. Health professionals will interact with the program as they might with a human consultant, describing the patient, receiving advice, and asking the consultant about the rationale for each recommendation. The initial prototype is using a knowledge base constructed by encoding the clinical expertise of a skilled psychiatrist in a set of rules. However, the researchers are identifying issues not well addressed by existing rule-based system-building tools (such as EMYCIN) and are anticipating considerable new research in the development of novel techniques for handling such problems.

MENTOR

The MENTOR project is directed by Dr. Stuart M. Speedie and Dr. Terrence F. Blaschke. Dr. Blaschke is Chief of the Division of Clinical Pharmacology in Stanford's Department of Medicine, and Dr. Speedie is a visiting scientist with the Division.

The goal of the MENTOR (Medical EvaluatioN of Therapeutic ORders) project is to design and develop an expert system for monitoring drug therapy for hospitalized patients that will provide appropriate advice to physicians concerning the existence and management of adverse drug reactions. The computer as a recording-keeping device is becoming increasingly common in hospital-based health care, but much of its potential remains unrealized. Furthermore, this information is provided to the physician in the form of raw data which is often difficult to interpret. The wealth of raw data may effectively hide important information about the patient from the physician. This is particularly true with respect to adverse reactions to drugs which can only be detected by simultaneous examinations of several different types of data including drug data, laboratory tests and clinical signs.

I.B.5. Major Books on Medical Artificial Intelligence

Just as the well known ***Handbook of Artificial Intelligence*** was developed on SUMEX several years ago, the resource has served as the focus for the development of two new books that are being published in 1984. Each book describes research projects that were largely dependent upon the SUMEX-AIM network for their successful implementation. Bruce Buchanan and Ted Shortliffe have edited a large collection of papers regarding the MYCIN system and its derivatives. They have also written new material and analyzed the results of the decade's experiments. The resulting volume, titled ***Rule-Based Expert Systems: The MYCIN Experiments of the Stanford Heuristic Programming Project***, will be published by Addison-Wesley in June.

A second volume, to be published by Addison-Wesley in July, is a collection of papers on AIM research efforts. The book, entitled ***Readings in Medical Artificial Intelligence: the First Decade***, was edited by Bill Clancey and Ted Shortliffe. Its 21 chapters summarize much of the research that SUMEX has helped spawn.

I.B.6. Training in Medical Information Science

Stanford's nascent program in Medical Information Sciences, mentioned briefly in last year's annual report, has matured significantly in the past 12 months. There will be 9 trainees in the program in September 1984, 7 working towards PhD degrees and 2 towards the MS degree. Of these trainees, 7 have MD degrees or are concurrently enrolled as medical students. Two of the trainees are playing central roles in the PATHFINDER research mentioned above, and several others are involved in ongoing AIM research using SUMEX facilities. The program has been awarded post-doctoral training support from the National Library of Medicine, received an equipment gift of four 9836 workstations from Hewlett Packard Company, and has received additional industrial and foundation grants for student support. We believe that SUMEX has been an important element in the rich medical computing research environment at Stanford that has in turn led to the successful implementation of this novel training effort. It is our belief that the medical computing and AIM communities, as well as biomedicine in general, will benefit greatly from an increased number of people trained to undertake research at the interface between medicine and computer science.

I.C. Administrative Changes

Carole Miller, who had served as the Administrative Assistant for SUMEX since 1974, accepted a new position as the Administrative Assistant of the Heuristic Programming Project in August of 1983. Carole has since moved on to become the Administrative Services Manager for the Center for Research on International Studies here at Stanford.

Patricia (Patti) M. McCabe has succeeded Carole as the Administrative Assistant for SUMEX-AIM. Patti comes to SUMEX-AIM from the Sponsored Projects Office at Stanford University where she was responsible for contracts and grant management, and was the primary liaison between Stanford University and the National Institutes of Health.

Roy Maffly stepped down as the SUMEX-AIM Liaison to devote more time to his responsibilities within the Stanford Medical Center. Larry Fagan, who returned to Stanford this past year as a Senior Research Associate in the Department of Medicine, has taken over for Roy as the new SUMEX AIM liaison.

I.D. Resource Management and Allocation

The mission of SUMEX-AIM, locally and nationally, entails both the recruitment of appropriate research projects interested in medical AI applications and the catalysts of interactions among these groups and the broader medical community. These user projects are separately-funded and autonomous in their management. They are selected for access to SUMEX on the basis of their computer and biomedical scientific merits, as well as their commitment to the community goals of SUMEX. Currently active projects span a broad range of applications areas such as clinical diagnostic consultation, molecular biochemistry, molecular genetics, medical decision making, and instrument data interpretation (Descriptions of the individual collaborative projects are in Section II beginning on page 69).

I.D.1. Management Committees

Since the SUMEX-AIM project is a multilateral undertaking by its very nature, several management committees have been created to assist in administering the various portions of the SUMEX resource. As defined in the SUMEX-AIM management plan adopted at the time the initial resource grant was awarded, the available facility capacity is allocated 40% to Stanford Medical School projects, 40% to national projects, and 20% to common system development and related functions. Within the Stanford aliquot, Prof. Feigenbaum and the BRP have established an advisory committee to assist in selecting and allocating resources among projects appropriate to the SUMEX mission. The current membership of this committee is listed in Appendix A.

For the national community, two committees serve complementary functions. An *Executive Committee* oversee's the operations of the resource as related to national users and renders final decisions on authorizing admission for new projects and revalidating continued access for existing projects. It also establishes policies for resource allocation and approves plans for resource development and augmentation within the national portion of SUMEX (e.g., hardware upgrades, significant new development projects, etc.). The Executive Committee oversees the planning and implementation of the AIM Workshop series, and assures coordination with other AIM activities as well. The Committee will continue to play a key role in assessing the possible need for additional future AIM community computing resources and in deciding the optimal placement and management of such facilities. The current membership of the Executive Committee is listed in Appendix A.

The Executive Committee met in 1983 during the AIM Workshop and via teleconferencing sessions. Items addressed during the committee meetings were final decisions on admissions of new AIM pilot projects, and the annual re-evaluation of continued access for AIM projects. In the latter area, a decision was reached after long and careful review to phase the SECS project out of SUMEX-AIM. The committee was concerned over the system impact of this project versus the current relevance and innovativeness of its research for AI. The implementation of this decision will be to phaseout SECS use in a fair and orderly manner, allowing for reduced system use until the completion of existing project commitments in March, 1985.

Reporting to the Executive Committee, an ***Advisory Group*** represents the interests of medical and computer science research relevant to AIM goals. The Advisory Group serves several functions in advising the Executive Committee: 1) recruiting appropriate medical/computer science projects, 2) reviewing and recommending priorities for allocation of resource capacity to specific projects based on scientific quality and medical relevance, and 3) recommending policies and development goals for the resource. The current Advisory Group membership is given in Appendix A.

These committees have actively functioned in support of the resource. Except for meetings held during the AIM workshops, the committees have "met" by messages, net-mail, and telephone conference, owing to the size of the groups and to save the time and expense of personal travel to meet face-to-face. The telephone meetings, in conjunction with terminal access to related text materials, have served quite well in accomplishing the agenda business. Other solicitations of advice requiring review of sizeable written proposals are done by mail.

We will continue to work with the management committees to recruit the additional high-quality projects which can be accommodated and to evolve resource allocation policies which appropriately reflect assigned priorities and project needs. We will continue to make information available about the various projects both inside and outside of the community and thereby promote the kinds of exchanges exemplified earlier and made possible by network facilities.

I.D.2. New Project Recruiting

We continue to see a very strong interest in Artificial Intelligence applications to medicine. We receive several inquiries a week, stimulated by information on SUMEX-AIM or the SUMEX-AIM subprojects. We are actively recruiting the best of these inquiries as pilot projects to provide new activities to replace projects that have matured and moved off of the SUMEX-AIM machine. A presentation was made at the American Association of Artificial Intelligence conference in August, 1983 to provide general information about SUMEX-AIM and encourage additional users. Additional information about SUMEX-AIM projects is available through well-attended presentations at national conferences in Artificial Intelligence. In addition, interest in the Artificial Intelligence approach to medical decision making has strongly increased in the national medical computing conferences. SUMEX-AIM related researchers are often the key personnel at these presentations.

During the Fall of 1983, two national and two Stanford-related projects were initiated. Many other interested researchers took advantage of SUMEX's ability to allow experimental access to existing computer programs. In addition, some of the more stable software for developing medical applications is now provided on tape for implementation on host computers outside of the SUMEX-AIM environment.

The criteria for the acceptance of new pilot projects continues to concentrate on the potential for excellence, and the novelty of the proposed concepts. We continue to seek projects that will extend our understanding of basic science issues underlying the application of the artificial intelligence approach to medical decision making. Thus, a project that will break new ground will be preferred to a project that uses existing ideas in a new area of medicine. We also encourage pilot projects to collaborate with of the existing bases of expertise in artificial intelligence techniques. Developing a new pilot project now requires more background and understanding of previous work in AI in medicine. However, the time needed to build a first prototype version may be substantially decreased by the use of packages developed by other SUMEX-AIM projects. SUMEX-AIM provides a unique opportunity for the development of pilot projects. We hope to build the number of pilot projects consistent with SUMEX resources and the availability of worthy project proposals.

I.D.3. Stanford Community Building

The Stanford community has undertaken several internal efforts to encourage interactions and sharing between the projects centered here. Professor Feigenbaum organized a project with the goal of assembling a handbook of current and state-of-the-art AI concepts and techniques. This project has had enthusiastic support from the students, and the work has culminated in the publication of a three-volume handbook set named the Handbook of AI, published by William Kaufman Press.

Weekly informal lunch meetings (SIGLUNCH) also are held between community members to discuss general AI topics, concerns and progress of individual projects, or system problems as appropriate. In addition, presentations are invited from a substantial number of outside speakers.

I.D.4. Existing Project Reviews

We have conducted a continuing careful review of on-going SUMEX-AIM projects to maintain a high scientific quality and relevance to our medical AI goals and to maximize the resources available for newly-developing applications projects. At meetings of the AIM Advisory Group and Executive Committee this past year, all of the national AIM projects were reviewed. These groups recommended continued access for most formal projects on the system, and the phaseout of the SECS project, details of which are covered on page 62.

I.D.5. Resource Allocation Policies

Policies have been established to control the allocation of critical facility resources (file space and central processor time) on the SUMEX-AIM 2060. File space management begins with an allocation of file storage, defined for each authorized project in consultation with the management committees. This allocation for any given project is redistributed among project members as directed by the individual principal investigators. System enforcement of project allocations is done on a weekly basis. As the weekly file dump is done, if the aggregate space in use by a project exceeds its allocation, files are archived from associated user directories which are over allocation until the project is within its authorized limits.

We are using the TOPS-20 class scheduler to attempt to enforce the 40:40:20 balance in terms of CPU utilization and to avoid system and user inefficiencies under overload conditions. In practice, the 40:40 split between Stanford and non-Stanford projects is fairly well realized (see Figure 10 on page 34 and the tables of recent project usage on page 36).

Our job-scheduling controls bias the allocation of CPU time based on per cent time consumed relative to the time allocated according to the 40:40:20 community split. However, the controls are "soft" in that they do not waste computer cycles if users below their allocated percentages are not on the system to consume those cycles. In the early years, the operating disparity in CPU use reflected a substantial difference in demand between the Stanford community and the developing national projects, rather than inequity of access. For example, the Stanford utilization is spread over a large part of the 24-hour cycle, while national-AIM users tend to be more sensitive to local prime-time constraints. (The 3-hour time zone phase shift across the continent is of substantial help in load-balancing). During peak times under the overload control system reported previously, the Stanford community experienced mutual contentions and delays while the AIM group had relatively open access to the system.

This disparity in usage has disappeared in recent years with the growth of the national user community, and we enabled overload controls for the national community as well. For the present, we propose to continue our policy of "soft" allocation enforcement for the fair split of resource capacity.

Our system also categorizes users in terms of access privileges. These comprise fully-authorized users, pilot projects, associates, guests, and network visitors in descending order of system capabilities. We want to encourage bona fide medical and health research people to experiment with the various programs available with a minimum of red tape, while not allowing unauthenticated users to bypass the advisory group screening procedures by coming on as guests. So far, we have had relatively little abuse compared to that experienced by other network sites, perhaps because of the personal attention directed by senior staff to logon records, and to other security measures. However, the experience of most other computer managers behooves us to be cautious about being as wide open as might be preferred for informal service to pilot efforts and demonstrations. We will continue developing this mechanism in conjunction with management committee policy decisions.

We also have encouraged mature projects to apply for their own machine resources in order to preserve the SUMEX-AIM resource for research and development efforts and to support projects unable to justify their own machines. The Rutgers resource has its own 2060 machine, part of which is allocated for AIM use, and the CADUCEUS project has installed a VAX 11/780 machine to support its planned development and program testing work. Profs. Lesgold and Greeno's "Simulation of Cognitive Processes" Project has moved entirely to their own local VAX.

I.E. Dissemination Efforts

Throughout its existence, SUMEX-AIM has expended substantial effort toward disseminating information about its activities as a resource and about the work of individual collaborative projects. We continue to make many presentations at professional meetings, to provide services to demonstrate developed AI programs to interested groups and individuals, to welcome visitors, and to work in organizing workshops within the SUMEX-AIM community to introduce our research to collaborating professional communities. We also directed considerable effort in the past toward working with the Research Resources Information Center to produce the "Seeds of Artificial Intelligence" monograph and other publications and press articles to address a broader community of technical and lay people.

Software Distribution

SUMEX continues to support various projects in the distribution of versions of their software to requesting individuals or groups. Following is a summary of software dissemination this past year:

EMYCIN — Both the "executable" and "source" versions of the EMYCIN distribution package were restructured for clarity and ease of installation. Thirty copies of the EMYCIN package have been generated for distribution of which about 6 were sources only. An Interlisp-VAX version of EMYCIN is now available, thanks to Ray Bates of USC-ISI, who did the conversion. This runs under UNIX and VMS.

AGE — Twenty-two copies of the AGE system have been distributed. Nearly half of these have been copies requested in ANSI format indicating they were evidently going to non-Tops20 sites (probably Vaxes). As with the EMYCIN system, Ray Bates at USC-ISI has converted AGE to run under Interlisp-VAX. A version is also available for the Xerox 1108 series Lisp workstations.

GENET — In conjunction with the phaseout of the GENET community on SUMEX, a software package comprised of programs and databases developed by researchers at Stanford and elsewhere was assembled for distribution to interested GENET users. Versions of the software were provided for use on both DEC-10 and DEC-20 systems operating under TOPS-10, TENEX, and TOPS-20. Installation procedures were documented, and a substantial amount of telephone consultation was provided. The package has been well-received and appears to be in active use at many of the 21 academic sites to which it was sent. Only one copy of the complete Genet system was set out in the past year. However, several sets of Genet related data files have been distributed. This includes several copies of the NIH and EMBL Sequence Libraries. A limited amount of operations support has been given to Brutlag's interaction with Sam Karlin of the Math department and a variety of other groups.

MRS — Twenty-two copies of MRS have been distributed through Sumex. Several others have also been distributed directly by the HPP. Most have been sent out to VAX/Unix sites or Symbolics Lisp machine sites.

SACON	Two copies of SACON have been prepared and distributed.
GLISP	Two copies of GLISP were distributed.

I.F. Comments on the Biotechnology Resources Program

Resource Organization

We continue to believe that the Biotechnology Resources Program is one of the most effective vehicles for developing and disseminating technological tools for biomedical research. The goals and methods of the program are well-designed to encourage building of the necessary multi-disciplinary groups and merging of appropriate technological and medical disciplines. In our experience with the SUMEX-AIM resource, several elements of this approach seem to emerge as key to the development and management of an effective resource:

1. Effective Management Framework -- There needs to be an explicit agreement between the BRP and the resource principal investigator which establishes a clear mandate for the resource and its allocation, provides worthwhile incentives for the host institution and investigator to invest the necessary substantial professional career time to develop and manage the resource, and ensures equitable distribution of resource services to its target community.

2. Close Working Relationship with the NIH -- A resource is a major and often long-term investment of money and human energy. A close and mutually-supportive working relationship between resource management, its advisory committees, and the NIH administration is essential to assure healthy development of the resource and its relationship to its user community. We at SUMEX-AIM have benefited immensely from such a relationship with Dr. William R. Baker, Jr., in the evolution of the SUMEX-AIM community. We look forward to a continuing mutually beneficial relationship with Dr. Baker's successor at the NIH.

3. Freedom to Explore Resource Potential -- A resource, by its nature, operates at the "cutting edge" in developing its characteristic technology and learning to effectively disseminate it to the biomedical community at large. The BRP should not impose artificial constraints on the resource for commercializing its efforts (fees for service) or developing its potential (funding duration limits or annual budget ceilings). Such artificial policy impositions can serve to undermine the very goals central to the BRP's reason for existence. Satisfactory policies in this regard have been worked out and should be retained.

Electronic Communications

SUMEX-AIM has pioneered in developing more effective methods for facilitating scientific communication. Whereas face-to-face contacts continue to play a key role, in the longer-term we feel that computer-based communications will become increasingly important to the NIH and the biomedical community. We would like to see the BRP take a more active role in promoting these tools within the NIH and its grantee community.

II. Description of Scientific Subprojects

II.A. Scientific Subprojects

The following subsections report on the AIM community of projects and "pilot" efforts including local and national users of the SUMEX-AIM facility at Stanford. However, those projects admitted to the National AIM community which use the Rutgers-AIM resource as their home base are not explicitly reported here.

In addition to these detailed progress reports, abstracts for each project and its individual users are submitted on a separate Scientific Subproject Form. However, we have included here briefer summary abstracts of the fully-authorized projects in Appendix B on page 209.

The collaborative project reports and comments are the result of a solicitation for contributions sent to each of the project Principal Investigators requesting the following information:

I. SUMMARY OF RESEARCH PROGRAM
- A. Project rationale
- B. Medical relevance and collaboration
- C. Highlights of research progress
  - --Accomplishments this past year
  - --Research in progress
- D. List of relevant publications
- E. Funding support

II. INTERACTIONS WITH THE SUMEX-AIM RESOURCE
- A. Medical collaborations and program dissemination via SUMEX
- B. Sharing and interactions with other SUMEX-AIM projects (via computing facilities, workshops, personal contacts, etc.)
- C. Critique of resource management (community facilitation, computer services, communications services, capacity, etc.)

III. RESEARCH PLANS
- A. Project goals and plans
  - --Near-term
  - --Long-range
- B. Justification and requirements for continued SUMEX use
- C. Needs and plans for other computing resources beyond SUMEX-AIM
- D. Recommendations for future community and resource development

We believe that the reports of the individual projects speak for themselves as rationales for participation. In any case, the reports are recorded as submitted and are the responsibility of the indicated project leaders. The only exceptions are the respective lists of relevant publications which have been uniformly formatted for parallel reporting on the Scientific Subproject Form.

II.A.1. Stanford Projects

The following group of projects is formally approved for access to the Stanford aliquot of the SUMEX-AIM resource. Their access is based on review by the Stanford Advisory Group and approval by Professor Feigenbaum as Principal Investigator.

In addition to the progress reports presented here, abstracts for each project and its individual users are submitted on a separate Scientific Subproject Form.

II.A.1.1. EXPEX - Expert Explanation Project

EXPEX - Expert Explanation Project

Edward H. Shortliffe, M.D., Ph.D.
Departments of Medicine and Computer Science
Stanford University

I. SUMMARY OF RESEARCH PROGRAM

A. Project Rationale

EXPEX is not a single project but a combination of efforts that are directed at basic issues in the development of representational schemes to facilitate knowledge acquisition and explanation. The work includes not only the study of fundamental representational formalisms but also the encoding of various types of knowledge, such as causal information and user models. In addition, to complement these research directions, the project has served as the focus for preparing three books on medical computing research.

We believe that the productivity of basic computer science research tends to be heightened by experiments that deal with significant real world problem domains. Challenges drawn from chemistry, medicine, and molecular biology have introduced additional complexity to expert systems work at Stanford, but have simultaneously forced system developers to respond to pragmatic constraints and user demands that have had a significant impact on the basic AI techniques selected or developed. Thus, we believe that creative investigation into symbolic reasoning techniques is facilitated by working in real world settings where the application forces us to avoid oversimplification. Much of our research effort therefore deals with medical domains (viz., endocrinology and renal pathophysiology).

B. Medical Relevance and Collaboration

Our interest in explanation derives from the insights we gained in developing explanatory capabilities for the MYCIN system. In the case of MYCIN and its descendents, we have been able to generate intelligible explanations by taking advantage of its rule-based representation scheme. Rules can be translated into English for display to a user, and their interactions can also be explicitly demonstrated. By adding mechanisms for understanding questions expressed in simple English, we were able to create an interactive system that allowed physicians to convince themselves that they agreed with the basis for the program's recommendations. The limitations of the explanations generated in this way have become increasingly obvious, however, and have led to improved characterization of the kinds of explanation capabilities that must be developed if clinical consultation systems are to be accepted by physicians. The potential use of workstation graphics as a means of avoiding natural language issues in the explanation process is also an area of great promise with which we are currently experimenting.

With these motivations in mind, we are involved in a series of research projects that address medical knowledge representation and explanation. The individual projects include the following:

1. Mr. Greg Cooper's NESTOR program uses a detailed knowledge base regarding pathophysiologic relationships in hypercalcemia. The program is designed to critique a physician's hypothesis regarding a proposed explanation for a set of patient manifestations when an elevated serum calcium has been observed. Of particular interest is the techniques Cooper has developed for using knowledge of causality to avoid the assumption of conditional independence commonly used in Bayesian diagnosis systems.

2. Mr. John Kunz has represented the knowledge of renal pathophysiology, including the ***quantitative*** relationships that characterize the way in which the body manages water and electrolytes, to develop a consultation and analysis system (AI/MM) that melds mathematics and AI techniques.

3. Building on his earlier experience with developing an explanation capability for NEOMYCIN (in collaboration with the GUIDON project members as outlined elsewhere in this report), Dr. Glenn Rennels has begun to work on a new system that uses knowledge of medicine to help formulate and resolve complex decision analyses. Convinced that decision analytic techniques would be better accepted in medicine if the physician were to interact with a knowledge-based interface (rather than with the decision trees themselves), Dr. Rennels has made use of "influence diagrams" as a central method for guiding the interaction. The explanation issues become especially evident when an analysis is complete and his system needs to generate a defense for the recommendation it has made.

4. Mr. Curt Langlotz has continued to work on a hypothesis assessment module for the ONCOCIN system. This program uses a ***critiquing model*** which inherently involves advanced explanation techniques. The work uses the Xerox 1108 professional workstation (Dandelion) and is further described in the ONCOCIN Project portion of this annual report.

5. During 1983, Ms. Shoko Tsuji completed a project using the Xerox workstation to experiment with graphical techniques for examining, manipulating, expanding, and editing a large medical knowledge base. Also working in the context of ONCOCIN, her code was designed for use by knowledge engineers. The work has inspired subsequent work in building an interface for the non-programmer clinician who wishes to write and test new protocols in the ONCOCIN environment. The project is described in greater detail in the ONCOCIN portion of the annual report.

To complement these basic research activities, we have prepared two books on Artificial Intelligence in Medicine and are beginning work on a third (see Section C for details).

C. Highlights of Research Progress

C.1 The NESTOR System

NESTOR is intended to allow a user to input patient data plus a hypothesis, and then have the system critique that hypothesis in light of the data. The system, an evolving thesis project that is largely the work of Mr. Greg Cooper, relies on basic associational information drawn in part from the INTERNIST-I knowledge base but supplemented with causal and temporal associations.

The motivation behind this research is the conviction that physicians want active control of the diagnostic process and that they also want and need a system that explains, in a user-tailored way, its evaluation of the physician's hypothesis. There may be times when the user wants to give complete control to NESTOR and just be in a mode of answering questions, but we feel that this should be an option and not a requirement. It is observations such as these that have also accounted for the hypothesis assessment work underway in the ONCOCIN research, briefly mentioned above and further described in the section of this report dealing with that project.

The initial NESTOR system is now largely complete and is undergoing evaluation at this time. Of particular interest is the adequacy of the techniques developed for allowing NESTOR to avoid the traditional assumption of conditional independence used in Bayesian systems. Also, because NESTOR's probabilistic model is more formal than the ad hoc scheme used in, say, INTERNIST, the assumptions made by our system are more explicit.

We have also developed search techniques that allow NESTOR to explore efficiently a very large search space in order to find the most probable (multiple disease) hypothesis. This technique is general and can be applied to many nonmedical problems where the goal is to find the most probable hypothesis among many possibilities.

C.2 Integrating Mathematical Models with AI Methods

This research project, known as AI/MM, is the dissertation research of Mr. John Kunz. The system integrates AI and simple mathematics to analyze a physiological model. In a selected medical domain (renal physiology), we have built a computer program based on these techniques. It analyzes physiological behavior, diagnoses abnormality, and explains the rationale for its analyses. The program fits data to the model, identifies whether the data are abnormal, and identifies the possible causes and effects of any abnormalities. The physiological model is based on knowledge about anatomy, the behavior of the physiological system, and the mechanism of action of the system. It's validity has been tested by having it analyze many of the problems discussed in Valtin's text *Renal Function*.

The specific aims of this project have been to:

1. Develop a vocabulary for a physiological model. The vocabulary represents the "basic physiology" of a biological system and appears to be adequate to express the concepts included in an introductory professional-level physiology text.

2. Develop a reasoning system which can solve problems expressed in the vocabulary.

3. Demonstrate the basic necessity, appropriateness and limitations of the vocabulary and reasoning procedure.

C.3 Knowledge-Based Explanations in a Decision Analysis Environment

This new project, thesis research by Dr. Glenn Rennels, is motivated by the observation that AI techniques could greatly facilitate a user's effort to specify the details of a complex clinical decision task and to seek assistance with that task. Although decision theoretic notions have been shown to be relevant to such medical problems, they have largely been unused by clinicians, even when computer-based solutions have been offered. We believe that an intelligent system should be able to ***guide*** the definition of the decision task and ***explain*** the results of the analysis without requiring that a user be familiar with the underlying decision analytic techniques being used to solve the problem.

The basic notion is to use directed graphs, termed "influence diagrams" as a language for communication with a physician at a graphical display terminal. Nodes in these graphs are defined by the user who is seeking advice, and their structure and meaning is largely intuitive. The task of converting influence diagrams to decision trees is a knowledge-based problem that is potentially well-suited for a solution that uses AI methods. Similarly, the results of a decision analysis, including the sensitivity analysis, will need to be explained to the physician user in terms of influence diagrams and knowledge of the domain. The necessary knowledge structures are currently being designed, and an early prototype system is operational. The research uses a 9836 workstation donated to the Medical Information Sciences Training Program by Hewlett-Packard Company and soon to be networked to the SUMEX 2060.

C.4 Books on Medical Artificial Intelligence and Medical Computing

We have completed two books, both of which are in press and due to be published in mid-1984:

- Clancey, W.J. and Shortliffe, E.H. *Readings in Medical Artificial Intelligence: The First Decade.* Reading, MA: Addison-Wesley, 1984.
- Buchanan, W.J. and Shortliffe, E.H. *Rule-Based Expert Systems: the MYCIN Experiments of the Stanford Heuristic Programming Project.* Reading, MA: Addison-Wesley, 1984.

In addition, we have just begun work on a textbook for students beginning to study medical computing and artificial intelligence. This multi-authored volume should be completed in draft form by the end of 1984. A 1985 publication date is contemplated.

- Shortliffe, E.H., Wiederhold, G.C.M., and Fagan, L.M. *An Introduction to Medical Computer Science.* Reading, MA: Addison-Wesley (in preparation).

D. Publications Since January 1983

1. Shortliffe, E.H. Medical consultation systems: designing for doctors. In *Designing for Human-Computer Communication* (M.S. Sime and M.J. Coombs, eds.), Chapter 8, pp. 209-238, London: Academic Press, 1983.
2. Shortliffe, E. H. Medical Cybernetics: The Challenges of Clinical Computing. In *Technology International Stability, and Growth*, S. Basheer Ahmed and Alice P. Ahmed, editors; Chapter 12, pp. 148-165; Associated Faculty Press, Inc., Port Washington, New York, 1984.
3. (*) Shortliffe, E.H. and Fagan, L.M. Expert systems research: modeling the medical decision making process. In *An Integrated Approach to Monitoring* (J.S. Gravenstein, R.S. Newbower, A.K. Ream, and N.T. Smith, eds.), pp. 183-200, Woburn, MA: Butterworth's, 1983.
4. Duda, R.O. and Shortliffe, E.H. Expert systems research. *Science*, 220:261-268 (1983).
5. (*) Langlotz, C.P. and Shortliffe, E.H. Adapting a consultation system to critique user plans. *International Journal of Man-Machine Studies*, 19:479-496 (1983)
6. Shortliffe, E.H. Hypothesis generation in medical consultation systems: artificial intelligence approaches. In *MEDINFO 83* (J.H. van Bemmel, M. Ball, and O. Wigertz, eds.), pp. 480-483, North Holland, Amsterdam, 1983.

7. (*) Tsuji, S. and Shortliffe, E.H. Graphical access to the knowledge base of a medical consultation system. *Proceedings of AAMSI Congress 83*, pp. 551-555, San Francisco, Ca., May 1983.

8. Shortliffe, E.H. The science of biomedical computing. In *Meeting the Challenge: Informatics and Medical Education* (J.C. Pages, A.H. Levy, F. Gremy, and J. Anderson, eds.), pp. 1-10, Amsterdam: North-Holland, 1983. To be reprinted in Medical Informatics, 1984.

9. (*) Kunz, J.C., Shortliffe, E.H., Buchanan, B.G., and Feigenbaum, E.A. Comparison of techniques of computer-assisted decision making in medicine. In *Pure and Applied Biostructure* (Claudio Niccolini, Ed.), Singapore: World Press, 1983.

10. (*) Kunz, J.C., Shortliffe, E.H., Buchanan, B.G., Feigenbaum, E.A. Computer-assisted decision making in medicine. *Journal of Philosophy and Medicine*, Summer 1984 (in press).

11. (*) Hasling, D. W., Clancey, W. J., and Rennels, G. Strategic explanations for a diagnostic consultation system. *International Journal of Man-Machine Studies*, Spring 1984 (in press).

12. Shortliffe, E.H. Reasoning methods in medical consultation systems: artificial intelligence approaches (tutorial). *Computer Programs in Biomedicine*, January 1984 (in press).

E. Funding Support

Grant Title: "The Development of Representation Methods to Facilitate Knowledge Acquisition and Exposition in Expert Systems"
Principal Investigator: Edward H. Shortliffe
Agency: Office of Naval Research; ID Number: NR 049-479
Term: January 1981 to December 1983
Total award: $456,622

Grant Title: "Research on Introspective Systems"
Principal Investigator: Michael R. Genesereth
Agency: Office of Naval Research; ID Number: NR 049-479
Term: January 1984 to December 1986
Total award: $312,070

Grant Title: "Information Structure and Use in Knowledge-Based Expert Systems"
Principal Investigator: Bruce G. Buchanan
Agency: National Science Foundation; ID Number: 83-12148
Term: March 1984 to February 1987
Total award: $300,000 (includes indirect costs)

II. INTERACTIONS WITH THE SUMEX-AIM RESOURCE

A. Medical Collaborations and Program Dissemination via SUMEX

None of these new programs is yet ready for dissemination. They are mostly